Scientific Objectives

The twentieth century has been an era of striking progress in understanding the fundamental structure of matter, beginning with the discovery of quantum mechanics and atomic physics, progressing to nuclear physics, and culminating with the Standard Model of elementary particle physics. However, the traditional analytical tools to calculate atomic and nuclear systems have proven inadequate to extract many of the predictions of quantum chromodynamics (QCD), the component of the Standard Model describing the strong interactions. Our understanding of nature will remain fundamentally deficient until we know how the rich and complex structure of strongly interacting matter, which comprises most of the known mass of the universe, arises from the interactions among quarks and gluons.

At present, the only method by which one can determine the predictions of QCD from first principles, with controlled systematic errors, is through large scale numerical simulations. Over the last several years the development and refinement of numerical algorithms, coupled with major increases in the capabilities of massively parallel computers have brought QCD simulations to a new level. It is now possible to calculate a few crucial quantities to an accuracy comparable with their experimental determination. The strong coupling constant and the masses of the c and b quarks are notable examples. The age of terascale computing brings us to a new threshold enabling similar accuracy for an increasingly broad range of fundamental quantities. Terascale lattice calculations will be essential to understanding the structure of nucleons and other hadrons, to determining the properties of hadronic matter under extreme conditions, and to searching for new physics via precise calculations of the parameters of the Standard Model.

As a new generation of experimental facilities is poised to explore the next frontiers of nuclear and particle physics, there is a special urgency in applying terascale computers to the study of QCD. Major investments are being made in the B-factory at the Stanford Linear Accelerator Center (SLAC), the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory, and the Jefferson National Laboratory Continuous Electron Beam Accelerator Facility (CEBAF), the three newest experimental facilities in high energy and nuclear physics. The incremental cost to perform the required terascale calculations is a small fraction of the total cost of these facilities, but is absolutely essential to their scientific missions.

Our overall objectives will require a sustained multi-year research effort supported by the significant increases in computing capabilities envisioned for the SSI. However, the work can be broken down into milestones, each of which will represent a major step forward for the field. The major source of uncertainty in most current QCD calculations is the neglect of quark-antiquark creation and annihilation, that is virtual quark loops, in generating lattice configurations. Therefore, a first step in our SSI research will be to generate a library of lattices central to the study of the topics enumerated above.

Our first major physics project will involve hadronic matrix elements, knowledge of which is essential for full use of the DOE's experimental facilities. In particular, the three quantities B_K (related to $K-\bar{K}$ mixing), $f_B\sqrt{B_B}$ (related to $B-\bar{B}$ mixing) and $\xi = f_{B_s}\sqrt{B_{B_s}}/f_B\sqrt{B_B}$ (related to the ratio of mixing for strange and non-strange B-mesons), are calculated well enough now that we can predict the improvements possible with given increases in computational resources. All three are needed to constrain the parameters of the Standard Model. The ultimate aim is to reduce the theoretical error to a level at which errors in experimental inputs dominate, which may require theoretical errors as small as 5%. Present theoretical errors are estimated to be about 20% (10% for ξ).

A major uncertainty in the three quantities results from the lack of a reliable continuum extrapolation in the presence of virtual quark loops. Existing calculations including quark loops involve only a limited range of lattice spacings, $a \geq 0.1\,\mathrm{fm}$. A machine sustaining 200 Gflops would, in the timescale of a year, extend such simulations down to $a \geq 0.05\,\mathrm{fm}$, allowing a much more thorough study of the continuum extrapolation. Another important source of uncertainty is the omission of dynamical strange quarks. At relatively little cost, the dependence on the dynamical strange quark mass could be investigated for $a \sim 0.1\,\mathrm{fm}$. We estimate that the combined effect of these improvements would be to roughly halve the present errors in B_K, $f_B\sqrt{B_B}$, and ξ. In addition, the errors would be largely calculated rather than estimated. These improvements would have a major impact on the determinations of standard model parameters.

To reach our ultimate goal of 5% errors in these three quantities requires improved numerical control over chiral extrapolations to light quark masses. We estimate that one must use non-degenerate light sea quarks with masses down to $\sim m_s/8$ (m_s being the physical strange quark mass). To attain this milestone will require a machine sustaining two TeraFlops.

As indicated above, the configurations used in the study of weak decays will also be used in a wide variety of other projects, including precise determinations of the strong coupling constant, quark masses, hadron spectrum, form factors, and moments of structure functions. Although separate configurations will be required for the study of high temperature QCD, the spectrum studies will be vital for determining the energy scale.

Appropriateness for and Linkage to SSI

Our research is ideally matched to the goals of the SSI. The results will be crucial for several of the most important components of the DOE mission, yet it could not be carried out on less powerful supercomputers than those of the SSI. The need stems from structures that extend over four space-time dimensions, encompassing $O(10^9)$ variables, from the requirement of high statistics in stochastic simulations, and from the iterative application of CPU intensive algorithms, such as the solution of huge systems of sparse linear equations. The planned investigations will challenge the power of the SSI resources, but there is every reason to expect that, through the collaboration of lattice gauge theorists, computer scientists and applied mathematicians, it will be possible to tune systems and application software to an outstanding level of performance. Lattice gauge theorists have extensive experience as early users of parallel machines. Indeed, we have been among the first users of every major parallel platform, and have generally obtained outstanding performance. We are confident that we can provide an early showcase, with strategies for high performance of real value to other SSI participants.

In two decades of applied computational research, lattice gauge theorists have put considerable effort into algorithm development. Indeed the precision already achieved in the calculation of several important observables is due as much to the introduction of powerful new algorithms as to advances in hardware. In this project, theoretical physicists will collaborate with applied mathematicians and computer scientists in an aggressive program of algorithmic research. Recent exciting developments in computational particle theory demand the introduction of new algorithms for exploiting their full potential. These algorithms should find applications in other areas as well, and thus the project will entail important benefits for the whole community of scientists working on the SSI. In particular, our interest in sparse linear algebra solvers is shared by most, if not all, of the SSI community, while our work on algorithms for the simulation of many-fermion systems has overlap with that of material scientists and nuclear physicists. At the same time, we expect our work to directly benefit from the algorithmic and software advances created by other SSI applications and CSET projects.

Organization of the Collaboration

The collaboration will consist of the X faculty members from Y universities and the W staff members from three national laboratories listed in Appendix A. It will be directed by a Steering Committee, consisting initially of Norman Christ, Michael Creutz, Paul Mackenzie, John Negele, Claudio Rebbi, Stephen Sharpe, and Robert Sugar, with Sugar serving as chair. The Steering Committee will assume overall responsibility for assuring that the goals of the collaboration are achieved. It will assign specific physics or computing projects to subgroups of the collaboration, and will, in consultation with the relevant physics subgroups, decide on the parameters for the dynamical fermion lattice sets which will be generated. Prospective new collaboration members who can make a significant contribution to the goals of the collaboration will be accepted upon a majority vote of the existing members.

Budget Justification

Our estimated budget for the first year of this proposal is $2,860,000. The three main components are for the support of scientific personnel, collaborative interactions and equipment. We propose to support five postdoctoral research associates, ten graduate students, three software engineers, a machine operator, a senior computer scientist and a senior applied mathematician. The latter two are important for expanding collaborative work on algorithms and software development. We also propose funds to support two group meetings per year for our widely distributed collaboration, and to purchase tools which will enable video teleconferencing via the Internet.

The single largest item in our budget is for the purchase of a prototype machine, which we envision having a peak speed of 200 to 400 Gflops. It would enable us to carry out code development prior to the installation of the first major SSI computer, so that we will be in a position to test, benchmark and begin substantive research as soon as it is installed. Even after the main SSI facility is in operation, a prototype machine will be very valuable for algorithm development, code optimization, and intermediate sized simulations which form an essential part of any large lattice gauge theory project. For these reasons, we place a very high priority on a prototype machine. Finally, we have included funds for a multi-processor computer to handle the numerous data analysis problems that are too large for a workstation and too small for the SSI or prototype machines.